Innov. Graph Theory **2**, 2025, pp. 157–190
https://doi.org/10.5802/igt.9

THE COMPLEXITY OF RECOGNIZING GEOMETRIC HYPERGRAPHS

by Daniel BERTSCHINGER, Nicolas EL MAALOULY, Linda KLEIST, Tillmann MILTZOW & Simon WEBER (*)

Abstract. — As set systems, hypergraphs are omnipresent and have various representations ranging from Euler and Venn diagrams to contact representations. In a *geometric representation* of a hypergraph $H = (V, E)$, each vertex $v \in V$ is associated with a point $p_v \in \mathbb{R}^d$ and each hyperedge $e \in E$ is associated with a connected set $s_e \subset \mathbb{R}^d$ such that $\{p_v \mid v \in V\} \cap s_e = \{p_v \mid v \in e\}$ for all $e \in E$. We say that a given hypergraph H is *representable* by some (infinite) family $\mathcal{F}$ of sets in $\mathbb{R}^d$, if there exist $P \subset \mathbb{R}^d$ and $S \subseteq \mathcal{F}$ such that (P, S) is a geometric representation of H. For a family $\mathcal{F}$, we define Recognition($\mathcal{F}$) as the problem to determine if a given hypergraph is representable by $\mathcal{F}$. It is known that the Recognition problem is $\exists\mathbb{R}$-hard for halfspaces in $\mathbb{R}^d$. We study the families of translates and homothets of balls and ellipsoids in $\mathbb{R}^d$, as well as of other convex sets, and show that their Recognition problems are also $\exists\mathbb{R}$-complete. In particular, for a *bi-curved, computable* set C, the recognition problem of the family of translates (or homothets) of C is $\exists\mathbb{R}$-complete if it is *T-difference-separable (H-difference-separable)*. We show that for bounded sets in the plane, convexity is equivalent to T-difference-separability and H-difference-separability ; in higher dimensions, convexity is necessary but not sufficient. Our results imply that these recognition problems are equivalent to deciding whether a multivariate system of polynomial equations with integer coefficients has a real solution.

1. Introduction

As set systems, hypergraphs appear naturally in various contexts, such as databases, clustering, and machine learning. They are also known as

Keywords: Hypergraph, geometric hypergraph, recognition, computational complexity, convex, ball, ellipsoid, halfplane, halfspace, translate, homothet, bi-curved, difference-separable.

2020 *Mathematics Subject Classification:* 05C62, 68R10, 05C65.

(*) A preliminary version appears in the Proceedings of GD 2023 [9].

We thank Patrick Schnider and Luis Montejano for their guidance towards references concerning the characterizations of ellipsoids.

T. M. is generously supported by the Netherlands Organisation for Scientific Research (NWO) under project no. VI.Vidi.213.150.

range spaces (in computational geometry) or *voting games* (in social choice theory). A hypergraph can be represented in numerous ways, e.g., by a bipartite incidence graph, a simplicial representation (if the set system is closed under taking subsets), Euler or Venn diagrams. Similar as in classic graph drawing, one can represent vertices by points and hyperedges by connected sets in $\mathbb{R}^d$ such that each set contains exactly the points of a hyperedge. For the purposes of legibility, uniformity, or also for aesthetic reasons, it is desirable that these sets satisfy additional properties, e.g., being convex or having similar appearance such as being homothetic copies or even translates of each other.

For an introductory example, suppose we are organizing a conference and have a list of accepted talks. Clearly, each participant wants to quickly identify talks of their specific interest. In order to create a good overview, we seek a way to nicely visualize this data. To this end, we label each talk by several keywords, e.g., `hypergraphs`, `complexity theory`, `planar graphs`, `beyond planarity`, `straight-line drawing`, `crossing numbers`, etc. Then, we create a representation, where each keyword is represented by a unit disk (or another nice geometric object of our choice) containing exactly the points that represent the talks that have this keyword. For an example of such a representation see Figure 1.1. In other words, we are interested in a geometric representation of the hypergraph where the vertex set is given by the talks and keywords define the hyperedges.

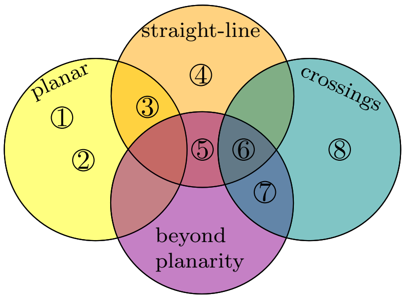

Figure 1.1. A geometric representation with unit disks of the hypergraph $H = (V, E)$ with $V = [8]$ and $E = \{\{1,2,3\},\{3,4,5,6\},\{5,6,7\},\{6,7,8\}\}$.

In this work, we investigate the complexity of deciding whether a given hypergraph has such a geometric representation. We start with a formal definition.

The Problem. In a *geometric representation* of a hypergraph $H = (V, E)$, each vertex $v \in V$ is associated with a point $p_v \in \mathbb{R}^d$ and each hyperedge $e \in E$ is associated with a connected set $s_e \subset \mathbb{R}^d$ such that $s_e \cap \{p_v \mid v \in V\} = \{p_v \mid v \in e\}$ for all $e \in E$. We say that a given hypergraph H is *representable* by some (possibly infinite) family $\mathcal{F}$ of sets in $\mathbb{R}^d$, if there exist a point set $P \subset \mathbb{R}^d$ and $S \subseteq \mathcal{F}$ such that (P, S) is a geometric representation of H. For a family $\mathcal{F}$ of geometric objects in $\mathbb{R}^d$, we define RECOGNITION($\mathcal{F}$) as the problem to determine whether a given hypergraph is representable by $\mathcal{F}$. Next, we give some definitions describing the geometric families studied in this work.

Bi-curved, Difference-separable, and Computable Sets. We study sets that are bi-curved, difference-separable and computable. While the first two properties are needed for $\exists\mathbb{R}$-hardness, the last one guarantees $\exists\mathbb{R}$-membership.

Let $C \subset \mathbb{R}^d$ be a set. We call C *computable* if for any point $p \in \mathbb{R}^d$ we can decide in polynomial time on a real RAM whether p is contained in C.

We say that C is *bi-curved* if there exists a unit vector $v \in \mathbb{R}^d$ and a neighborhood N of v, such that for every unit vector $v' \in N$ there exists a unique pair of tangent hyperplanes on C with normal vector v' with the following two properties: (i) C lies between the hyperplanes and (ii) each of these hyperplanes intersects the boundary of C in a single point. We furthermore require that the contact points change continuously when changing $v' \in N$. Note that a bi-curved set is necessarily bounded. As an example, each strictly convex bounded set in any dimension is bi-curved; for such a set, any unit vector v fulfills the conditions. As illustrated by Figure 1.2a, being strictly convex is not necessary for a set to be bi-curved.

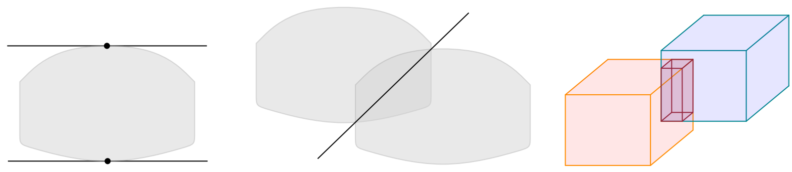

(a) This burger-like set is bi-curved as shown by the two tangent hyperplanes.

(b) A hyperplane separating the symmetric difference of two translates of the burger-like set.

(c) Two cubes in $\mathbb{R}^3$ whose symmetric difference cannot be separated by a plane.

Figure 1.2. Illustration for the notions bi-curved and difference-separable.

We call a family $\mathcal{F}$ of sets *difference-separable* if for any two members C_1, C_2 of $\mathcal{F}$, there exists a hyperplane which strictly separates $C_1 \setminus C_2$ from $C_2 \setminus C_1$. We call a set $C \subset \mathbb{R}^d$ *T-difference-separable* if the family of all translates of C in $\mathbb{R}^d$ is difference-separable. As we prove in Lemma 6.1, any bounded T-difference-separable set is necessarily convex. Consequently, bi-curved and T-difference-separable sets are always convex. Reversely, any bounded convex set in $\mathbb{R}^2$ is T-difference-separable, see Figure 1.2b for an example. Thus, for bounded sets in $\mathbb{R}^2$, convexity and T-difference-separability are equivalent. However, in higher dimensions this is not the case: There exists a convex set in $\mathbb{R}^d$ such that the family of its translates is not difference-separable, e.g., consider the two cubes in $\mathbb{R}^3$ in Figure 1.2c. Besides families of translates, we are interested in families of homothets and call a set C H-difference-separable if the family of homothets of C is difference-separable. Interestingly, any convex set in $\mathbb{R}^2$ and every ellipsoid is also H-difference-separable; for proofs of these facts we refer to to Lemmata 6.3 and 6.4. In conclusion, the bi-curved and H-difference-separable families include all strictly convex sets in $\mathbb{R}^2$ and families of balls and ellipsoids are H-difference-separable in all dimensions. We are not aware of other natural geometric families with those two properties in all dimensions. Note that balls and ellipsoids are naturally computable.

We are now ready to state our results.

1.1. Results

We study the recognition problem of geometric hypergraphs. We first consider the maybe simplest type of geometric hypergraphs, namely those that stem from halfspaces. It is known due to Tanenbaum, Goodrich, and Scheinerman [64] that the Recognition problem for geometric hypergraphs of halfspaces is NP-hard, but their proof actually implies $\exists\mathbb{R}$-hardness as well. We present a slightly different proof of this fact due to two reasons. Firstly, their proof lacks details about extensions to higher dimensions. Secondly, it is a good stepping stone towards Theorem 1.3.

Theorem 1.1 (Tanenbaum, Goodrich, Scheinerman [64]). — *For every $d \geqslant 2$, Recognition($\mathcal{F}$) is $\exists\mathbb{R}$-complete for the family $\mathcal{F}$ of halfspaces in $\mathbb{R}^d$.*

We note that for $d = 1$, the recognition problem for halfspaces can be solved in polynomial time, because it is easy to decide whether two hyperedges are represented by halfspaces that are unbounded into the same or different directions.

Next we consider families of objects that are translates or homothets of a given object. For $d = 1$, the considered RECOGNITION problems can be solved in polynomial time as the problems are very close to recognizing unit interval graphs and interval graphs, respectively.

LEMMA 1.2. — *Let $C \subseteq \mathbb{R}$ be a convex set, and let $\mathcal{T}_C$ be the family of all translates of C and $\mathcal{H}_C$ be the family of all homothets of C. Then RECOGNITION($\mathcal{T}_C$) and RECOGNITION($\mathcal{H}_C$) are polynomial time solvable.*

For $d \geqslant 2$, we show $\exists\mathbb{R}$-completeness.

THEOREM 1.3. — *theorem For $d \geqslant 2$, let $C \subseteq \mathbb{R}^d$ be a bi-curved, T-difference-separable (thus convex) and computable set, and let $\mathcal{T}_C$ be the family of all translates of C. Then RECOGNITION($\mathcal{T}_C$) is $\exists\mathbb{R}$-complete.*

With similar techniques, we determine the recognition complexity for families of homothetic objects.

THEOREM 1.4. — *For $d \geqslant 2$, let $C \subseteq \mathbb{R}^d$ be a bi-curved, H-difference-separable (thus convex) and computable set, and let $\mathcal{H}_C$ be the family of all homothets of C. Then RECOGNITION($\mathcal{H}_C$) is $\exists\mathbb{R}$-complete.*

Together with our insights on T-difference-separable sets, namely Lemmata 6.3 and 6.4, Theorems 1.3 and 1.4 yield the following result.

COROLLARY 1.5. — *When C is a bounded strictly convex set in $\mathbb{R}^2$, a bi-curved convex set in $\mathbb{R}^2$, or an ellipsoid in $\mathbb{R}^d, d \geqslant 2$, then the decision problems RECOGNITION($\mathcal{T}_C$) and RECOGNITION($\mathcal{H}_C$) are $\exists\mathbb{R}$-hard.*

It is natural to wonder which sets are T-difference-separable and H-difference-separable. In the plane, we obtain the following simple characterization.

THEOREM 1.6. — *The following statements are equivalent for any bounded set C in $\mathbb{R}^2$:*

(1) C is convex.
(2) C is T-difference-separable.
(3) C is H-difference-separable.

However, as mentioned before, in higher dimensions convexity is not a sufficient criterion for a set to be T-difference-separable (or even H-difference-separable). As it turns out compact H-difference-separable sets in $\mathbb{R}^3$ are exactly the ellipsoids.

THEOREM 1.7. — *A compact set C in $\mathbb{R}^3$ is H-difference-separable if and only if it is an ellipsoid.*

One might be under the impression that the Recognition problem is $\exists\mathbb{R}$-complete for every reasonable family of geometric objects of dimension at least two. However, we show that the problem is contained in NP for translates as well as for homothets of polygons and thus, if $\mathsf{NP} \subsetneq \exists\mathbb{R}$, also not $\exists\mathbb{R}$-complete.

Theorem 1.8. — *Let P be a simple polygon with integer coordinates in $\mathbb{R}^2$. For the family $\mathcal{T}_P$ of all translates of P, Recognition($\mathcal{T}_P$) is contained in NP. Similarly, for the family $\mathcal{H}_P$ of all homothets of P, Recognition($\mathcal{H}_P$) is contained in NP.*

Organization. We discuss membership results in Section 2, i.e., we prove the membership parts of Theorems 1.1, 1.3 and 1.4 as well as Theorem 1.8 and Lemma 1.2. We introduce the version of pseudohyperplane stretchability used in our hardness reductions in Section 3. Proofs of the hardness parts of Theorems 1.1, 1.3 and 1.4 can be found in Sections 4 and 5, respectively. In Section 6, we present present the characterizations of difference-separable sets, namely we prove Theorems 1.6 and 1.7. We conclude with interesting future directions in Section 7.

1.2. Related Work

In this section, we present an overview over related work on the complexity class $\exists\mathbb{R}$, geometric intersection graphs, and on other set systems related to hypergraphs.

The Existential Theory of the Reals. The complexity class $\exists\mathbb{R}$ (pronounced as 'ER' or 'exists R') is defined via its canonical complete problem ETR (short for *Existential Theory of the Reals*) and contains all problems that polynomial-time many-one reduce to it. In an ETR instance, we are given a sentence of the form

$$\exists x_1, \ldots, x_n \in \mathbb{R} : \varphi(x_1, \ldots, x_n),$$

where φ is a well-formed and quantifier-free formula consisting of polynomial equations and inequalities in the variables and the logical connectives $\{\wedge, \vee, \neg\}$. The goal is to decide whether this sentence is true.

The complexity class $\exists\mathbb{R}$ gains its importance from its numerous influential complete problems. Important $\exists\mathbb{R}$-completeness results include the realizability of abstract order types [44, 59], geometric linkages [51], and the recognition of geometric intersection graphs, as further discussed below.

More results concern graph drawing [22, 23, 36, 52], the Hausdorff distance [31], polytopes [21, 49], Nash-equilibria [8, 11, 12, 27, 55], training neural networks [3, 10], matrix factorization [19, 56, 57, 58, 65], continuous constraint satisfaction problems [43], geometric packing [5], the art gallery problem [2, 63], and covering polygons with convex polygons [1]. For more details, consider the survey by Schaefer, Cardinal, and Miltzow [53].

Geometric Hypergraphs. Many aspects of hypergraphs with geometric representations have been studied. Hypergraphs represented by touching polygons in $\mathbb{R}^3$ have been studied by Evans et al. [25]. Bounds on the number of hyperedges in hypergraphs representable by homothets of a fixed convex set have been established by Axenovich and Ueckerdt [7]. Smorodinsky studied the chromatic number and the complexity of coloring of hypergraphs represented by various types of sets in the plane [61]. Dey and Pach [20] generalize many extremal properties of geometric graphs to hypergraphs where the hyperedges are induced simplices of some point set in $\mathbb{R}^d$. Haussler and Welzl [29] defined ϵ-nets, subsets of vertices of hypergraphs called range spaces with nice properties. Such ϵ-nets of geometric hypergraphs have been studied quite intensely [6, 39, 46, 47].

While there are many structural results, we are not aware of any research into the complexity of recognizing hypergraphs given by geometric representations, other than the recognition of embeddability of simplicial complexes, as we will discuss in the next paragraph.

Other Representations of Hypergraphs. Hypergraphs are in close relation with abstract simplicial complexes. In particular, an abstract simplicial complex (complex for short) is a set system that is closed under taking subsets. A k-complex is a complex in which the maximum size of a set is k. In a geometric representation of an abstract simplicial complex $H = (V, E)$ each ℓ-set of E is represented by a ℓ-simplex such that two simplices of any two sets intersect exactly in the simplex defined by their intersection (and are disjoint in case of an empty intersection). Note that 1-complexes are graphs and hence deciding the representability in the plane corresponds to graph planarity (which is in P). In stark contrast, Abrahamsen, Kleist and Miltzow recently showed that deciding whether a 2-complex has a geometric embedding in $\mathbb{R}^3$ is $\exists\mathbb{R}$-complete [4]; they also prove hardness for other dimensions. Similarly, piecewise linear embeddings of simplicial complexes have been studied [15, 16, 17, 38, 40, 42, 60].

Recognizing Geometric Intersection Graphs. Given a set of geometric objects, its intersection graph has a vertex for each object, and an edge between any two intersecting objects. The complexity of recognizing geometric intersection graphs has been studied for various geometric objects. We summarize these results in Figure 1.3.

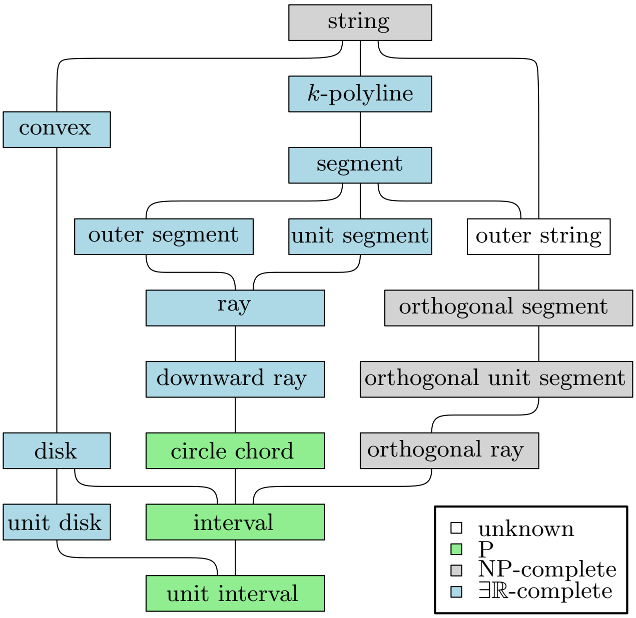

Figure 1.3. Containment relations of geometric intersection graphs. Recognition of a green class is in P, of a grey class is NP-complete, of a blue class is ∃ℝ-complete, and of a white class is unknown.

While intersection graphs of circle chords (Spinnrad [62]), unit intervals (Looges and Olariu [35]) and intervals (Booth and Lueker [13]) can be recognized in polynomial time, recognizing string graphs (Schaefer and Sedgwick [54]) is NP-complete. In contrast, ∃ℝ-completeness of recognizing intersection graphs has been proved for (unit) disks by McDiarmid and Müller [41], convex sets by Schaefer [50], downward rays by Cardinal et al. [18], outer segments by Cardinal et al. [18], unit segments by Hoffmann et al. [30], segments by Kratochvíl and Matoušek [34], k-polylines by Hoffmann et al. [30], and unit balls by Kang and Müller [33].

The existing research landscape indicates that recognition problems of intersection graphs are ∃ℝ-complete in case that the family of objects satisfy two conditions: Firstly, they need to be "geometrically solid", i.e., not strings. Secondly, some non-linearity must be present by either allowing rotations, or by the objects having some curvature. Our results indicate that this general intuition might translate to the recognition of geometric hypergraphs.

2. Membership

In this section we show $\exists\mathbb{R}$-membership parts of Theorems 1.1 and 1.3, as well as Theorem 1.8 and Lemma 1.2.

2.1. Halfspaces

For a given hypergraph H, it is not difficult to formulate an ETR formula describing all needed properties for a geometric representation by halfspaces. Therefore, we get the $\exists\mathbb{R}$-membership part of Theorem 1.1.

LEMMA 2.1. — *Fix $d \geqslant 1$ and let $\mathcal{F}$ denote the family of halfspaces in $\mathbb{R}^d$. Then RECOGNITION($\mathcal{F}$) is contained in $\exists\mathbb{R}$.*

Proof. — For a given hypergraph H, we formulate an ETR formula as follows. For each vertex/point, we create variables $p = (p_1, \dots, p_d)$ to represent the point. Similarly, for each hyperedge/halfspace, we create variables $h = (h_1, \dots, h_{d+1})$ to represent the coefficients of the halfspace. Then for each point p that is supposed to be in some halfspace h, we create the constraint:

$$h_1 p_1 + \dots h_d p_d \leqslant h_{d+1}.$$

Similarly, if p is not contained in a halfspace h, we create the constraint:

$$h_1 p_1 + \dots h_d p_d > h_{d+1}.$$

This is a valid ETR sentence that is equivalent to the representability of H. Note that for any fixed dimension d the ETR sentence is of polynomial size. □

The $\exists\mathbb{R}$-membership part of Theorem 1.3 is obtained by providing a simple verification algorithm [24] (similar to how NP-membership can be shown), based on the fact that our considered set C is computable.

LEMMA 2.2. — *For some $d \geqslant 1$, let $C \subseteq \mathbb{R}^d$ be a computable set and let $\mathcal{T}_C$ be the family of all translates of C. Then, RECOGNITION($\mathcal{T}_C$) is contained in $\exists\mathbb{R}$.*

Recall that the class NP is usually described by the existence of a witness and a verification algorithm. The same characterization exists for $\exists\mathbb{R}$ using a real verification algorithm. Instead of the witness consisting of binary words of polynomial length, in addition a polynomial number of real-valued numbers are allowed as a witness. Furthermore, in order to be able to use those real numbers, the verification algorithm is allowed to work on the so-called real RAM model of computation. The real RAM allows arithmetic operations with real numbers in constant time [24].

Proof. — We describe a real verification algorithm as mentioned above. The witness consists of the (real) coordinates of the points representing the vertices and the coefficients of the translation vectors representing the hyperedges. By definition of computable, a verification algorithm can efficiently check if each point is contained in the correct sets. □

2.2. Translates and Homothets of Polygons

Here, we show Theorem 1.8, i.e., NP-membership of Recognition of translates of a simple polygon P.

Theorem 1.8. — *Let P be a simple polygon with integer coordinates in $\mathbb{R}^2$. For the family $\mathcal{T}_P$ of all translates of P, Recognition($\mathcal{T}_P$) is contained in* NP*. Similarly, for the family $\mathcal{H}_P$ of all homothets of P, Recognition($\mathcal{H}_P$) is contained in* NP*.*

Proof. — The proof uses a similar argument to the one used to show that the problem of packing translates of polygons inside a polygon is in NP [5]. We wish to find a certificate that can be tested in polynomial time using linear programming, where the variables of the linear program correspond to the translation and scaling vectors of our homothets, as well as the coordinates of our points. To do this, the certificate needs to specify linear constraints for both containment and non-containment of a point p in a homothet of P. To get such constraints, we first triangulate the convex hull of P, such that each edge of P appears in the triangulation. Then, a representation of a hypergraph H by homothets of P gives rise to a certificate as follows: For each pair of a point p and a homothet P' of P, we specify whether p lies in the convex hull of P'. If it does, we specify in which triangle p lies. Otherwise, we specify an edge of the convex hull for which p and P' lie on opposite sides of the line through the edge. For an illustration, consider Figure 2.1.

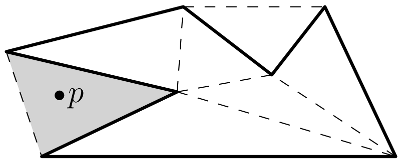

Figure 2.1. The polygon P', a triangulation of its convex hull, and the triangle that contains p.

Such a certificate can be tested in polynomial time: we create a linear program whose variables describe the locations of the points p and the translation vectors and scaling factors of each homothet of P, and whose constraints enforce the points to lie in the regions described by the certificate. This linear program has a number of constraints and variables polynomial in the size of H, and can thus be solved in polynomial time.

The solution of this linear program gives the location of the points and the translation vectors and scaling factors of the polygons. This implies that these coordinates are all polynomial and could be used as a certificate directly.

For RECOGNITION($\mathcal{T}_{\mathcal{P}}$), we use the same machinery but fix the scaling factors to be 1. □

2.3. 1D Versions

Next, we show that the one-dimensional problems can be solved in polynomial time.

LEMMA 1.2. — *Let $C \subseteq \mathbb{R}$ be a convex set, and let $\mathcal{T}_C$ be the family of all translates of C and $\mathcal{H}_C$ be the family of all homothets of C. Then* RECOGNITION($\mathcal{T}_C$) *and* RECOGNITION($\mathcal{H}_C$) *are polynomial time solvable.*

Proof. — We can assume that C is bounded because the problem is trivial for $C = \mathbb{R}$, and for C being a halfspace both RECOGNITION($\mathcal{T}_C$) and RECOGNITION($\mathcal{H}_C$) are solved by testing whether the vertices of H can be sorted such that each hyperedge forms a prefix, i.e., the hyperedges ordered by inclusion are a total order. It thus remains to consider the case where C is an interval, and thus all hyperedges have to be represented by intervals (homothets), or by unit intervals (translates). In the following, we will ensure a representation (if it exists) where all event points are unique. Thus, we may assume without loss of generality that C is closed.

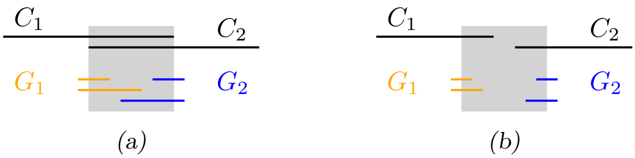

Figure 2.2. Illustration of the proof of Lemma 1.2: removing undesired intersections without introducing containment relations.

Let $H = (V, E)$ be a hypergraph. We construct the following incidence graph $G = (U, F)$. For each vertex $v \in V$, we add a vertex u_v to U and for each hyperedge $e \in E$, we add a vertex u_e to U. Moreover, we insert an edge $u_v u_e$ in F if $v \in e$ and an edge $u_e u_{e'}$ if $e \cap e' \neq \emptyset$. In the following we show how we can check properties of G in order to determine whether H is representable by $\mathcal{H}_C$ and/or $\mathcal{T}_C$.

Homothets. As a first step, we argue that H has a geometric representation with homothets of C if and only if G is an interval graph. We consider representations where the endpoints of intervals (and points) are unique, i.e., no two endpoints/points have the same coordinate.

Consider an interval representation of G. By construction of G, the neighborhood $N(u_v)$ of each u_v is a clique. By Helly's theorem, there thus exists a point p within the interval of u_v that intersects all intervals of $N(u_v)$. Hence, we may shrink the interval of each u_v to the point p. This yields a geometric representation of H with intervals, i.e., H is representable by $\mathcal{H}_C$.

The reverse direction is a little bit more intricate. Note that in a representation (P, S) of H, intervals (representing the hyperedges E) may intersect even if they do not share any point in P. However, we may (easily) modify R to remove these undesired intersections. To this end, consider two sets C_1 and C_2 with an undesired non-empty intersection $C_1 \cap C_2$. Without loss of generality, we may assume that $C_1 \cap C_2$ corresponds to the interval between the start of C_2 and the end of C_1; otherwise one set contains the other and contains no point from P. Now, we consider all event points within $C_1 \cap C_2$ and partition the corresponding sets into two groups G_1 and G_2 depending on whether they are end points or start points within $C_1 \cap C_2$, respectively; note that no set C_i is contained in $C_1 \cap C_2$ as it would contain no point from P. For an illustration, consider Figure 2.2a.

We reassign the event points within $C_1 \cap C_2$ such that all events of G_1 appear before G_2 and the order within each group is maintained, see Figure 2.2b. Note that this operation never changes whether or not some set contains another. (We remark that we could just swap the end of C_1 with the start of C_2 when just considering homothets, however this would not maintain the no-containment-property which we use in the case of translates below.) After removing all undesired properties, we have an interval representation of G.

Thus, when considering homothets, we may check in time $O(|U| + |F|)$ whether G is an interval graph and return the answer [14, 28].

Translates. Similarly, it holds true that H has a geometric representation with translates of C if and only if G has an interval representation where no hyperedge-interval contains another hyperedge-interval; we call such a interval representation *h-proper*. The only-if direction follows from the modifications of an interval representation represented above; we may clearly assume that G is an interval graph, otherwise H is not representable by $\mathcal{T}_C$. For the other direction, recall that a graph is a unit interval graph if and only if it is a proper interval graph, i.e., no interval contains another interval. In particular, from any proper interval representation, a unit interval representation with the same ordering of the start (and end vertices) can be computed in polynomial time [26, Section 3].

Consequently, as a first step, we check whether $G[E]$ has a proper interval representation. If not, we return no. Otherwise, we consider such a representation R and aim to incorporate points for V. To do so, we exploit the following fact. The left-right ordering of a twin-free unit interval graph is unique (up to reversing); we consider two vertices u and v of a graph to be *twins* if $N[v] = N[u]$, i.e., if their closed neighborhoods coincide. Hence the only degrees of freedom of R is the ordering between vertices with same neighborhoods (twins). Suppose that in R all intervals of twins coincide.

Consider a set S of twins and the union of points P_S contained in at least one but not all of them. For an illustration, consider the thick interval in Figure 2.3a. There is a natural poset on P_S where each point p is associ-

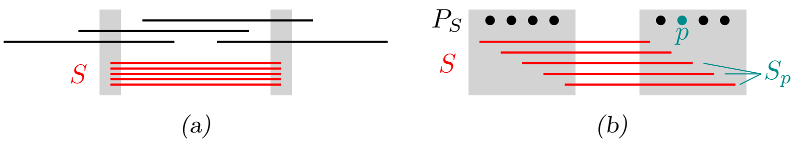

Figure 2.3. Illustration of the proof of Lemma 1.2 for translates. Considering a set S of twins (in red), we aim to find an ordering such that the points of P_S can be added as illustrated in (b).

ated with the subset S_p of twins containing it and the subsets are ordered by inclusion; without loss of generality we may assume that no two S_p's coincide, otherwise we treat the corresponding points momentarily as one point and insert more copies in small enough vicinity at the very end. It is easy to see that the width of this poset is at most 2, otherwise G is not h-proper. To this end, note that any representation restricted to $S \cup P_S$ is essentially as depicted in Figure 2.3b, where some of the depicted points may not be present. Clearly, the left group of points and the right group of points form a chain. Thus, the two chains of the poset allow us to partition P_S into two sets P_ℓ and P_r; if S is not an isolated clique, a neighbor will make the left/right groups unique, otherwise we call them arbitrarily. In order to place $p \in P_\ell$, the intervals of S_p must start before the intervals in $S \setminus S_p$. Similarly, to place $p \in P_r$, the intervals of S_p must end after the intervals in $S \setminus S_p$. If there exists a total order of the twins that satisfies all conditions, we insert them; otherwise we have a no-certificate. We repeat this process until no two intervals coincide.

Finally, we check if for each point there exists a suitable location. If so, we have a desired representation. Otherwise, by the uniqueness of twin-free representations and our careful insertion of twins, there exists no representation. □

We remark that the recognition problem of interval graphs where each vertex is represented by a point or a unit interval was studied by Rautenbach and Szwarcfiter [48]; however, in their setting each vertex may be represented by either object while the assignment is fixed in our setting.

3. Pseudohyperplane Stretchability

In this section we introduce the problem used to show the $\exists\mathbb{R}$-hardness parts of Theorems 1.3 and 1.4.

A *hyperplane arrangement* in $\mathbb{R}^d$ is an arrangement of hyperplanes in $\mathbb{R}^d$. We are interested in a generalization. A *pseudohyperplane arrangement* in $\mathbb{R}^d$ is an arrangement of *pseudohyperplanes*, where a pseudohyperplane is a set homeomorphic to a hyperplane, and each intersection of pseudohyperplanes is homeomorphic to a plane of some dimension. In the classical definition, every set of d pseudohyperplanes has a non-empty intersection. Here, we consider *partial pseudohyperplane arrangements (PPHAs)*, where not necessarily every set of $\leqslant d$ pseudohyperplanes has a common intersection.

A PPHA is *simple* if, for all $k = 1, \ldots, d$, no more than k pseudohyperplanes intersect in a space of dimension $d - k$. In particular, no $d + 1$

pseudohyperplanes have a common intersection. We call the 0-dimensional intersection points of d pseudohyperplanes the *vertices* of the arrangement. A simple PPHA $\mathcal{A}$ is *stretchable* if there exists a hyperplane arrangement $\mathcal{A}'$ such that each vertex in $\mathcal{A}$ also exists in $\mathcal{A}'$ and each (pseudo-)hyperplane splits this set of vertices the same way in $\mathcal{A}$ and $\mathcal{A}'$. In other words, each vertex of $\mathcal{A}$ lies on the correct side of each hyperplane in $\mathcal{A}'$. We then call the hyperplane arrangement $\mathcal{A}'$ a *stretching* of $\mathcal{A}$.

The decision problem d-Stretchability asks whether a simple PPHA in $\mathbb{R}^d$ is stretchable. For $d = 2$, d-Stretchability contains the stretchability of simple pseudoline arrangements which is known to be $\exists\mathbb{R}$-hard [44, 59]. It is straightforward to prove $\exists\mathbb{R}$-hardness for all $d \geqslant 2$.

Theorem 3.1. — *d-Stretchability is $\exists\mathbb{R}$-hard for all $d \geqslant 2$.*

Proof. — We reduce from stretchability of simple pseudoline arrangements, which is $\exists\mathbb{R}$-hard as shown in [44, 59].

Consider a simple pseudoline arrangement L in the x_1x_2-plane. We consider $d-2$ pairwise orthogonal hyperplanes $h^1, \ldots, h^{d-2}$ whose common intersection is the x_1x_2-plane; e.g., the hyperplanes defined $x_i = 0$ for $i = 3, \ldots, d$. The intersection of these hyperplanes serves as a canvas in which we aim to embed L. We extend each pseudoline ℓ to a pseudohyperplane h_ℓ by extending it orthogonally to all $h^1, \ldots, h^{d-2}$, see Figure 3.1.

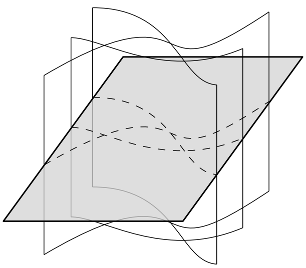

Figure 3.1. Extending a simple pseudoline arrangement (dashed) to a partial pseudohyperplane arrangement in $\mathbb{R}^3$. The grey hyperplane serves as the "canvas".

Clearly, the resulting pseudohyperplane arrangement $\mathcal{A}$ can be built in polynomial time. Note that all intersection points of d pseudohyperplanes in $\mathcal{A}$ correspond to intersection points of L.

If L is stretchable, $\mathcal{A}$ is clearly stretchable, as the above construction can be applied to the stretched line arrangement of L.

If $\mathcal{A}$ is stretchable, L is stretchable, since restricting each hyperplane h_ℓ to the intersection of the hyperplanes $h^1, \ldots, h^{d-2}$ yields a line arrangement which is equivalent to L.

As we have thus reduced stretchability of simple pseudoline arrangements to d-Stretchability, this concludes the proof. □

Similar extensions of pseudoline stretchability to higher dimensions have been studied in the literature. For example, Mnëv's universality theorem [44] extends to higher dimensions, however we are not aware of any existing proofs that it also implies $\exists\mathbb{R}$-hardness in $d > 2$. Kang and Müller [33] also studied a similar version of stretchability of partial arrangements of pseudohyperplanes.

4. Hardness for Families of Halfspaces

We now present the hardness part of Theorem 1.1.

Theorem 1.1 (Tanenbaum, Goodrich, Scheinerman [64]). — *For every $d \geqslant 2$, Recognition($\mathcal{F}$) is $\exists\mathbb{R}$-complete for the family $\mathcal{F}$ of halfspaces in $\mathbb{R}^d$.*

Proof. — We reduce from d-Stretchability. Let $\mathcal{A}$ be a simple PPHA. For an example consider Figure 4.1a. In a first step, we insert a parallel twin ℓ' for each pseudohyperplane ℓ. The twin is close enough to ℓ such that ℓ and ℓ' have the same intersection pattern. Since ℓ and ℓ' are parallel, they do not intersect each other. This yields an arrangement $\mathcal{A}'$, see Figure 4.1b.

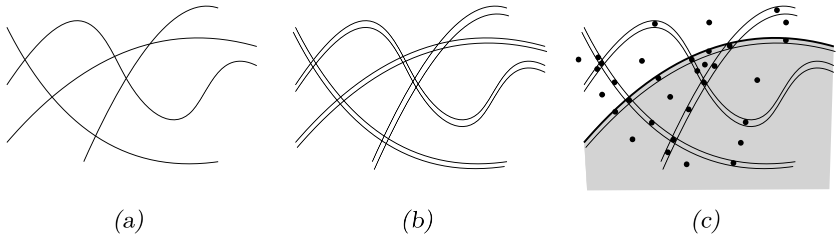

Figure 4.1. Construction of the hypergraph H. (a) A simple PPHA $\mathcal{A}$. (b) The arrangement $\mathcal{A}'$ obtained by inserting twins. (c) The vertices of H are the points in the cells, hyperedges of H are defined by the pseudohalfspaces; the gray region shows one of the hyperedges.

In a second step, we introduce a point in each d-dimensional cell of $\mathcal{A}'$; each point represents a vertex in our hypergraph H. Lastly, we define a hyperedge for each pseudohyperplane ℓ of $\mathcal{A}'$: The hyperedge contains all of the points that lie on the side of ℓ that the twin pseudohyperplane ℓ' lies in, see Figure 4.1c. Note that we define such a hyperedge for every pseudohyperplane of $\mathcal{A}'$. Thus, for every pseudohyperplane ℓ of the original arrangement $\mathcal{A}$ we define two hyperedges, whose union contains all vertices of H.

It remains to show that H is representable by halfspaces if and only if $\mathcal{A}$ is stretchable. If $\mathcal{A}$ is stretchable, the construction of a representation of H is straightforward: Consider a hyperplane arrangement $\mathcal{B}$ which is a stretching of $\mathcal{A}$. Then, for each hyperplane, we add a parallel hyperplane very close, so that their intersection patterns coincide. This results in a hyperplane arrangement $\mathcal{B}'$. We now prove that every d-dimensional cell of $\mathcal{A}'$ must also exist in $\mathcal{B}'$. We show this by considering the 0-dimensional cells, which we call vertices. First, note that each d-dimensional cell of $\mathcal{A}'$ corresponds to a cell of $\mathcal{A}$, which has at least one vertex on its boundary. All vertices of $\mathcal{A}$ exist in $\mathcal{B}$ by definition of a stretching. Furthermore, the subarrangement of the d hyperplanes in $\mathcal{B}$ intersecting in this vertex must be simple, since their intersection could not be 0-dimensional otherwise. In the twinned hyperplane arrangement $\mathcal{B}'$, all 3^d of the d-dimensional cells incident to the parallelotope formed by the planes through this vertex and their twinned copies (a cell is given by the following choice for each of the hyperplane pairs: above both hyperplanes, between the hyperplanes, or below both hyperplanes) must exist. This proves that all d-dimensional cells of $\mathcal{A}'$ also exist in $\mathcal{B}'$. Inserting a point in each such d-dimensional cell and considering the (correct) halfspaces bounded by the hyperplanes of $\mathcal{B}'$ yields a representation of H.

We now consider the reverse direction. Let $(P, \mathcal{H})$ be a tuple of points and halfspaces representing H. Let $h_{i,1}$ and $h_{i,2}$ be the two halfspaces associated with a pseudohyperplane ℓ_i of $\mathcal{A}$. Let p_i denote the $(d-1)$-dimensional hyperplane bounding $h_{i,1}$. We show that the family $\{p_i\}_i$ of these hyperplanes is a stretching of $\mathcal{A}$.

For each intersection point q of d pseudohyperplanes $\ell_1, \dots \ell_d$ in $\mathcal{A}$, we consider the corresponding $2d$ pseudohyperplanes in $\mathcal{A}'$. The PPHA $\mathcal{A}'$ contains 3^d d-dimensional cells incident to their 2^d intersections; each of which contains a point. We first show that the associated halfspaces must induce at least 3^d cells, one of which is bounded and represents the intersection point, see also Figure 4.2a: These 3^d points have pairwise distinct characteristics of whether or not they are contained in each of the $2d$ halfspaces , i.e., for every pair of points there exists a halfspace containing one but not the other. Thus, these points need to lie in distinct cells of the arrangement of halfspaces, which proves the claim.

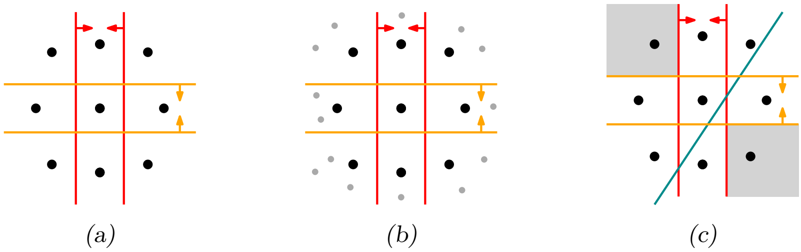

Figure 4.2. Illustration for the proof of Theorem 1.1 for $d = 2$ showing that representability of H implies stretchability of $\mathcal{A}$. (a) Any two pseudolines ℓ_i, ℓ_j in $\mathcal{A}$ have four corresponding lines bounding the respective halfplanes in H; these four lines induce 9 cells, each of which contains a point. (b) Each point in P belongs to exactly one of these 9 cells; the central bounded cell contains a unique point representing the intersection of ℓ_i and ℓ_j. (c) If the central bounded cell was intersected by a line p_k with $k \neq i, j$, then the 9 points do not lie on one side of p_k.

Moreover, every point in P belongs to exactly one of these 3^d cells. In particular, the central bounded cell, denoted by $c(q)$, contains exactly one point of P, see Figure 4.2b.

Now, we argue that the complete cell $c(q)$ (and thus in particular the intersection point of the hyperplanes representing q) lies on the correct side of each hyperplane p in $\{p_i\}_i$. Note that, by construction of the hypergraph H, the 3^d points of q lie on the same side of p. Suppose for a contradiction that p intersects $c(q)$, see Figure 4.2c. Then there exist two unbounded cells incident to $c(q)$ which lie on different sides of p; these cells can be identified by translating p until it intersects $c(q)$ only in the boundary. This yields a contradiction to the fact that the 3^d points of q lie on the same side of p.

We conclude that each intersection point of d pseudohyperplanes in $\mathcal{A}$ also exists in the arrangement $\{p_i\}_i$ and lies on the correct side of all hyperplanes. Thus, $\{p_i\}_i$ is a stretching of $\mathcal{A}$ and we conclude that $\mathcal{A}$ is stretchable. $\square$

5. Hardness for Families of Translates and Homothets

We are now going to prove the hardness part of Theorems 1.3 and 1.4. To this end, consider any fixed bi-curved, and T-difference-separable set C

in $\mathbb{R}^d$. Recall that thus C is convex. Note that we can assume C to be fully-dimensional, since otherwise each connected component would live in some lower-dimensional affine subspace, with no interaction between such components. We use the same reduction from the problem d-STRETCHABILITY as in the proof for halfspaces in the previous section, i.e., given a simple PPHA $\mathcal{A}$ we perform the doubling procedure and define the hypergraph H as illustrated in Figure 4.1. We will now show that this constructed hypergraph H is representable by translates of C if and only if the given PPHA $\mathcal{A}$ is stretchable. In fact, we present two lemmas, one for each direction, that are strong enough to imply both, Theorems 1.3 and 1.4.

LEMMA 5.1. — *If $\mathcal{A}$ is stretchable, H is representable by translates of C.*

The idea behind the proof is that a stretching $\mathcal{A}'$ of $\mathcal{A}$ can be scaled and stretched in such a way that every hyperplane has a normal vector close to the vector v witnessing that C is bi-curved, and such that all the vertices lie within some sufficiently small box. Then, for every halfspace $h_\pm$ bounded by some hyperplane h in $\mathcal{A}$', there exists a translate of C which approximates $h_\pm$ within the small box. This intuition is shown in Figure 5.1a. Since the hyperplane arrangement is simple, and there is some slack between the hyperplanes bounding the two twin halfspaces (as we argued above in the proof of Theorem 1.1), such an approximation is sufficient.

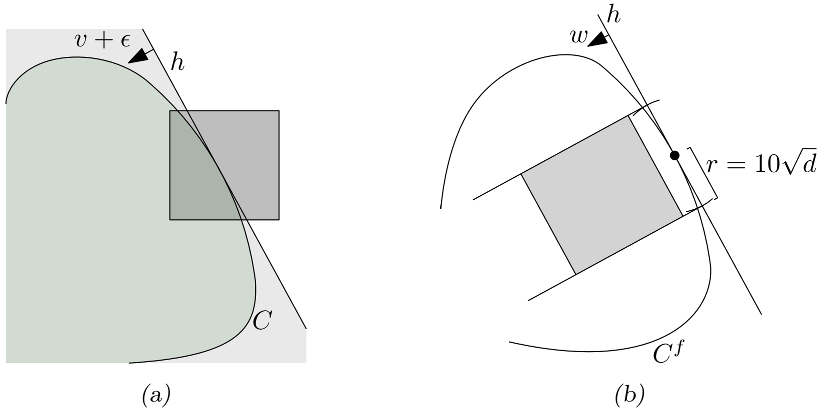

Figure 5.1. Illustrations for the proof of Lemma 5.1. a) Within the small box (dark grey), the translate of C (green) approximates the halfspace (light grey) bounded by h. b) An illustration of the requirement on the scaling factor f. The set C^f must contain the grey region.

Proof. — We assume that $\mathcal{A}$ is stretchable. We already proved in the previous section that thus there exists an arrangement of hyperplanes, in which we can create a twin of each hyperplane (with a tiny distance α between the twins), and in which we can place all the vertices of H in the appropriate d-dimensional cells. If a vertex is placed between two twin hyperplanes, we assume it to be equidistant to them. As before, we denote this arrangement of hyperplanes and points by $\mathcal{B}'$.

Let v be the unit vector certifying that C is bi-curved; recall the definition in Section 1. Therefore, there exists $\epsilon > 0$, such that any unit vector w with $\|w - v\|_2 \leqslant \varepsilon$ also fulfills the conditions to certify that C is bi-curved.

We now assume that $\mathcal{B}'$ fulfills the following properties:

(1) the normal vectors of all hyperplanes have distance at most ε to v or to $-v$

(2) every intersection point of d hyperplanes as well as every point representing a vertex of H, is contained in $[-1,1]^d$.

Both properties can be achieved by applying some affine transformation with positive determinant, thus preserving the combinatorial structure of $\mathcal{B}'$.

To represent the hyperedges of H, we will now use very large copies of C. Note that technically we are not allowed to scale C, but scaling C by a factor f is equivalent to scaling the arrangement by a factor $1/f$. Let C^f be the set C scaled by factor f.

In order to determine the necessary scaling factor f, we consider the curvature of C^f in all the points where the tangent hyperplanes of C^f with normal vector w for $\|w - v\|_2 \leqslant \varepsilon$ intersect C^f. In each such tangent hyperplane h with (unit) normal vector w, we draw a $(d-1)$-ball B of radius $10\sqrt{d}$ around the touching point $h \cap C^f$. Note that $10\sqrt{d}$ is larger than the length of any line segment contained in the box $[-1,1]^d$. Now, f has to be large enough such that C^f contains every point $p + w \cdot \lambda$, for $p \in B$ and $\alpha/10 \leqslant \lambda \leqslant 10\sqrt{d}$. This ensures that the boundary of C^f does not curve away from the tangent hyperplane too quickly, and that C^f is "thick". In other words, C^f locally behaves like an only very slightly curved halfspace. See Figure 5.1 for an illustration of this requirement on C^f.

We now replace each hyperplane h of the arrangement $\mathcal{B}'$ by a translate C^f_h of C^f, placed such that h is a tangent hyperplane of C^f_h, the single point $h \cap C^f_h$ lies within the box $[-1,1]^d$, and C^f_h lies completely to the side of h containing its twin hyperplane. It remains to prove that C^f_h contains exactly those points of $\mathcal{B}'$ which are on this side of h. Firstly, C^f_h cannot contain more points, since C^f_h is a subset of the halfspace delimited by h containing its twin hyperplane. Second, we claim that C^f_h contains all these points. To see this, note that within the box $[-1,1]^d$ containing all points, the boundary of C_h^f is close enough to h that it must contain all points between h and its twin, since these points are located equidistant to the two hyperplanes. Furthermore, all points on the other side of the twin hyperplane are also contained in C_h^f since within the box $[-1,1]^d$, the boundary $\delta(C_h^f)$ lies completely between h and its twin hyperplane. $\square$

Now we consider the reverse direction.

LEMMA 5.2. — *If the hypergraph H is representable by a difference-separable family $\mathcal{F}$, then $\mathcal{A}$ is stretchable.*

Proof. — Consider a representation of H with $\mathcal{F}$. By construction, the two members $C_{i,r}, C_{i,l}$ of $\mathcal{F}$ corresponding to the two hyperedges of each pseudohyperplane ℓ_i must intersect as they contain at least one common point. For each pseudohyperplane ℓ_i of $\mathcal{A}$, we consider some hyperplane p_i which separates $C_{i,r} \setminus C_{i,l}$ from $C_{i,l} \setminus C_{i,r}$. Such a hyperplane exists because $\mathcal{F}$ is difference-separable. Let $\mathcal{P} := \{p_i\}_i$ be the hyperplane arrangement consisting of all these separating hyperplanes. We aim to show that $\mathcal{P}$ is a stretching of $\mathcal{A}$.

To this end, consider d pseudohyperplanes $\ell_1, \ldots, \ell_d$ which intersect in $\mathcal{A}$. Figure 5.2 displays the case $d = 2$. Furthermore, consider one more pseudohyperplane ℓ', and let p', C'_r, C'_l denote the separator hyperplane and members of $\mathcal{F}$ corresponding to ℓ'. We show that the intersection $I_p := p_1 \cap \ldots \cap p_d$ is a single point which lies on the same side of p' as the point $I_\ell := \ell_1 \cap \ldots \cap \ell_d$ lies of ℓ'.

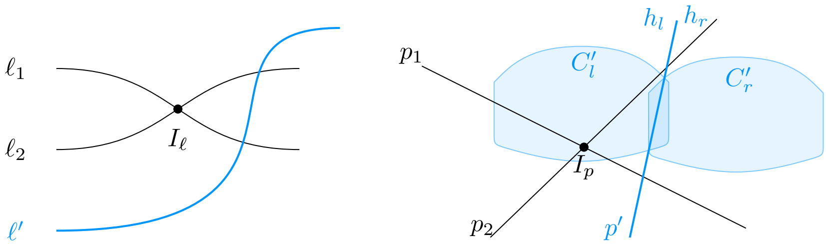

(a) Pseudohyperplanes ℓ_1, ℓ_2, ℓ' in $\mathcal{A}$ *(b)* Hyperplanes p_1, p_2, p' in $\mathcal{P}$.

Figure 5.2. Illustration for the proof of Lemma 5.2 for $d = 2$. Some pseudohyperplanes in $\mathcal{A}$ and their corresponding hyperplanes in $\mathcal{P}$.

The hyperplane p' divides the space into two halfspaces h_r and h_l such that $C'_r \backslash C'_l \subseteq h_r$ and $C'_l \backslash C'_r \subseteq h_l$. By construction, the two hyperedges defined for ℓ' cover all vertices of H and the vertices in the cells around I_ℓ

belong to only one hyperedge. Suppose without loss of generality that these vertices only belong to the hyperedge represented by C'_l. We will show that the intersection I_p must then be a point in h_l.

We first show that the intersection I_p is a point, i.e., 0-dimensional. Consider all 2^d d-dimensional cells of $\mathcal{A}$ around I_ℓ. The construction of H implies that each such cells contains a distinct point, and these points must all lie in distinct cells of the sub-arrangement of the involved hyperplanes $p_1, \dots, p_d$. Assuming that I_p is not a single point, this sub-arrangement is not simple, and the hyperplanes divide space into strictly fewer than 2^d cells, which results in a contradiction.

Next we prove that I_p is in h_l. Assume towards a contradiction that $I_p \in h_r$, see also Figure 5.3. Consider the d lines that are formed by the intersections of subsets of $d-1$ hyperplanes among $p_1, \dots, p_d$. Each of these lines is the union of two rays beginning at I_p. Observe that the hyperplane p' can only intersect one of the two rays forming each line. Let S be the convex cone centered at I_p defined by the d non-intersected rays. Observe that S does not intersect p', so S must be fully contained in h_r, i.e., $S \cap h_l = \emptyset$.

Note, however, by the construction of the hypergraph, there must be a point that lies in $S \cap (C'_l \setminus C'_r) \subseteq S \cap h_l$, which is a contradiction.

We conclude that $\mathcal{P}$ is a stretching of $\mathcal{A}$, and thus $\mathcal{A}$ is stretchable. □

Lemmas 5.1 and 5.2 combined now prove hardness of the recognition problem for the family of translates of C as well as for the superfamily of homothets. This is due to the fact that Lemma 5.1 guarantees a representation using only translates and Lemma 5.2 allows to reconstruct a line arrangement even from any representation with members of a difference-separable family. This completes the proof of Theorems 1.3 and 1.4.

Together with Lemmata 6.3 and 6.4, Theorems 1.3 and 1.4 imply Corollary 1.5.

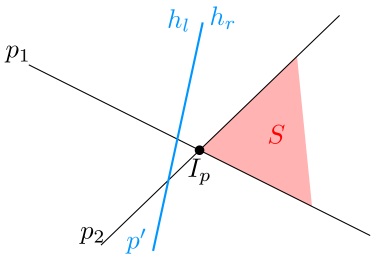

Figure 5.3. Illustration for the proof of Lemma 5.2 for $d = 2$. The cone S must intersect $C'_l \setminus C'_r$, which contradicts I_p lying in h_r.

6. Characterizations of Difference-Separable Sets

In this section, we aim to better understand which bounded sets are T-/H-difference-separable. As it turns out convexity is necessary. We show that in the plane convexity is also sufficient for H-difference-separability (and thus T-difference-separability), however this is not the case for higher dimensions. Afterwards, we characterize compact H-difference-separable sets in three dimensions.

6.1. In 2D

We present a full characterization of bounded sets in the plane, namely, we show that T-difference-separability and H-difference-separability are equivalent to convexity.

Theorem 1.6. — *The following statements are equivalent for any bounded set C in $\mathbb{R}^2$:*

(1) C is convex.
(2) C is T-difference-separable.
(3) C is H-difference-separable.

We show this statement in two steps. Firstly we show in Lemma 6.1 that convexity is necessary in all dimensions already for T-difference-separability. Secondly, we show that every convex bounded set in $\mathbb{R}^2$ is H-difference-separable.

Lemma 6.1. — *If $C \subset \mathbb{R}^d$ is bounded and T-difference-separable, then C is convex.*

Proof. — Consider a set C that is not convex. Then there exists a segment s such that its endpoints are contained in C and its midpoint is not. Without loss of generality we assume that s is vertical. We consider two translates C_1 and C_2 of C such that the midpoint of each segment coincides with an endpoint of the other segment, see also Figure 6.1. In other words, C_1 is shifted below C_2 by half the length of s.

Now, we consider the (vertical) line ℓ supporting s_1 (and s_2) and orient it from bottom to top. By boundedness and the fact that C_1 is shifted below C_2, ℓ intersects C_1 first and thus contains a point from $C_1 \setminus C_2$. Then ℓ meets the midpoint of s_1 which is contained in $C_2 \setminus C_1$, afterwards ℓ meets the midpoint of s_2 which is contained in $C_1 \setminus C_2$. As points from $C_1 \setminus C_2$ and $C_2 \setminus C_1$ lie alternatingly on a line, these two sets cannot be separated by a line. Thus C is not T-difference-separable. □

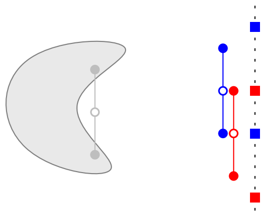

Figure 6.1. Illustration for the proof of Lemma 6.1 for $d = 2$.

We remark that with similar ideas we obtain a stronger necessary condition in general (without assuming boundedness): If C is T-difference-separable, then for every line ℓ, either $\ell \cap C$ is convex (if C is not bi-infinite in the direction of ℓ) or $\ell \cap \overline{C}$ is convex. It is easy to see that this condition is also sufficient in 1D. In particular, $C \subset \mathbb{R}^1$ is T-difference-separable if and only if C or its complement are convex (i.e., an interval).

Next we show that convexity implies H-difference-separability. To this end, we make use of the following fact (we present the statement in a more precise way).

Lemma 6.2 ([37], Corollary 2.1.2.2). — *The intersection of the boundaries of two different but homothetic convex compact sets in $\mathbb{R}^2$ consists of at most two connected components each of which is either a point or a segment.*

We make use of this fact in order to show the following.

Lemma 6.3. — *Every convex bounded set C in $\mathbb{R}^2$ is H-difference-separable.*

Proof. — Consider two homothetic copies C_1 and C_2 of C. If $C_1 \setminus C_2$ or $C_2 \setminus C_1$ is empty, we are done. Consequently, we assume that both difference sets are non-empty. By Lemma 6.2, the intersection of the boundaries ∂C_1 and ∂C_2 consists of at most two connected components each of which is either a point or a segment. As the statement is trivial for segments, we may assume that C is two-dimensional, i.e., the boundaries are closed Jordan curves. First, we present arguments for the case when C_1 and C_2 coincide on their boundary, i.e., if $(\partial C_1 \cap \partial C_2) \cap C_1 = (\partial C_1 \cap \partial C_2) \cap C_2$; e.g., this is the case if C is closed or open.

If $\partial C_1 \cap \partial C_2$ has one connected component, then, because C_1 does not contain C_2 and vice versa, either C_1 and C_2 only differ in their boundary or their interiors are disjoint; otherwise the Jordan curves intersect twice.

Hence, the hyperplane separation theorem for convex sets guarantees that C_1 and C_2 are difference-separable.

Now, we consider the case that $\partial C_1 \cap \partial C_2$ contains two components. In this case, we may choose any two points p_1 and p_2, one in each component, e.g., two closest points, see also Figure 6.2a. We prove that the line ℓ through p_1 and p_2 separates $C_1 \setminus C_2$ and $C_2 \setminus C_1$: Consider $x_1 \in C_1 \setminus C_2$ and suppose there exists $x_2 \in C_2/C_1$ on the same side of ℓ. Then, by the Jordan curve theorem and convexity, ∂C_1 and ∂C_2 intersect in a point $x \notin \ell$ such that the line supporting xp_i separates x_1 and x_2. By assumption, there exists an $i \in \{1, 2\}$ such that x and p_i lie in the same component of $\partial C_1 \cap \partial C_2$, i.e., both ∂C_1 and ∂C_2 contain the segment $p_i x$. However, because C_1 and C_2 do not differ in boundary points, this yields a contradiction to the fact that $x_i \in C_i \setminus C_{i\pm 1}$.

Figure 6.2. Illustration for the proof of Lemma 6.3.

It remains to consider the case that some boundary points of C belong to C while others do not. If $\partial C_1 \cap \partial C_2$ consists of a single segment s, then it remains to consider the case that the difference sets are contained in $\partial C_1 \cap \partial C_2$. By convexity, the points of $C_i \cap \partial C_i$ form an interval and the differences of two intervals can easily be separated.

If $\partial C_1 \cap \partial C_2$ consists of two segments s_1 and s_2, we may without loss of generality assume that the relative translation vector of C_1 and C_2 is (non-vanishing) horizontal, and C_1 is not right of C_2. It follows that s_1 and s_2 are horizontal; otherwise $\partial C_1 \cap \partial C_2$ consists of one component by convexity: Excluding potential horizontal segments, ∂C has a left and right boundary curve that starts from a topmost point and ends in a bottommost point. If $\partial C_1 \cap \partial C_2$ contains a non-horizontal segment, then this is supported by a segment from C_1 in the right boundary curve and a segment from C_2 in the left boundary curve. By convexity, the shared segments supports a separating line and thus C_1 and C_2 are otherwise disjoint.

Hence, on each segment s_j, we choose p_j as the leftmost point of s_j if $C_1 \cap s_j = \emptyset$, otherwise we choose the supremum (w.r.t. the x-coordinate) of $s_j \cap C_1$, i.e., the right boundary point of the segment $s_j \cap C_1$. We show that the line ℓ through p_1 and p_2 separates $C_1 \setminus C_2$ and $C_2 \setminus C_1$. Suppose there exists $x_i \in C_i \setminus C_{i\pm1}$ for $i \in \{1,2\}$ left of ℓ. We may assume that x_1 is an interior point of $C_1 \setminus C_2$ because C_2 is (2-dimensional and) a horizontal translate of C_1 lying to its right. As above, it follows that one x_i, namely x_2, belongs to $\partial C_1 \cap \partial C_2$, i.e., to some s_j. The fact that $x_2 \in s_j$ is left of p_j implies that $C_1 \cap s_j \neq \emptyset$. By definition of p_j, there exists a point in $s_j \cap C_1$ right of x_2. Moreover, as C_2 is a horizontal translate of C_1, there exists $x_2' \in C_1 \cap \ell_j$ left of x_2, see also Figure 6.2b. By convexity of C_1, it follows that $x_2 \in C_1$, a contradiction.

For the case that there exists $x_i \in C_i \setminus C_{i\pm1}$ for $i \in \{1,2\}$ right of ℓ, we may analogously assume that x_2 is an interior point. It follows that x_1 is contained in some s_j and right of p_j, a contradiction to the choice of p_j. □

Together, Lemmata 6.1 and 6.3 imply Theorem 1.6, so this finishes the proof.

6.2. In 3D

In three dimensions, we can show that H-difference-separable compact sets are exactly the ellipsoids.

Theorem 1.7. — *A compact set C in $\mathbb{R}^3$ is H-difference-separable if and only if it is an ellipsoid.*

We suspect this to be true in higher dimensions as well, however our proof only generalizes to higher dimensions for the following, simpler direction.

Lemma 6.4. — *Every ellipsoid in $\mathbb{R}^3$ is H-difference-separable (and thus also T-difference-separable).*

Proof. — As a set of homothetic ellipsoids can be transformed to a set of balls by an affine transformation, it suffices to show that a ball in $\mathbb{R}^d$ is H-difference-separable: It is a well-known fact that the intersection of two d-spheres in $\mathbb{R}^d$ is a $(d-1)$-sphere (or a point or empty) that lies in a hyperplane (orthogonal to the line connecting the centers of the spheres). For instance, consider a sphere S_1 of radius R centered at the origin and a sphere S_2 of radius r at $(\delta, 0, \ldots, 0)$ with $\delta > 0$ and $R \geqslant r$. If $\delta > R+r$, the intersection is empty. If $\delta = R + r$, the intersection is a point. Hence, we consider the case that $0 < \delta < r + R$. Any point in $S_1 \cap S_2$ must lie in the hyperplane $x_1 = \alpha$ for $\alpha := (r^2 - R^2 + \delta^2)/2\delta$. It is easy to check that the hyperplane $x_1 = \alpha$ separates $B_1 \setminus B_2$ from $B_2 \setminus B_1$: If $x \in B_1$ and $x_1 > \alpha$, then $x \in B_2$ because $(x_1 - \delta)^2 - x_1^2 \leqslant R^2 - r^2 \iff x_1 \geqslant \alpha$. Similarly, $x \in B_2$ and $x_1 < \alpha$ implies $x \in B_1$. □

Now, we consider the reverse direction. We show that H-difference-separability of C implies that every two parallel (nonempty) sections of C are homothetic. A *section* of C is the non-empty intersection of C with a plane. As a matter of fact this property characterizes ellipsoids in $\mathbb{R}^3$.

Theorem 6.5 (Kakeya [32] and Nakagawa [45]). — *Let $C \subset \mathbb{R}^3$ be a compact convex body. Every two parallel (non-empty) sections of C are homothetic if and only if C is an ellipsoid.*

It hence remains to prove the property of parallel sections. To this end, we start with a useful property of non-homothetic sets in the plane.

Lemma 6.6. — *Let A and B be two non-homothetic convex compact sets in $\mathbb{R}^2$. Then there exists a homothet A' of A and a homothet B' of B such that $B' \setminus A'$ and $A' \setminus B'$ cannot be separated by a line.*

Proof. — We first scale and translate A and B such that their areas and centroids coincide. We then argue that the resulting sets A' and B' fulfill the condition. For an illustration, consider Figure 6.3.

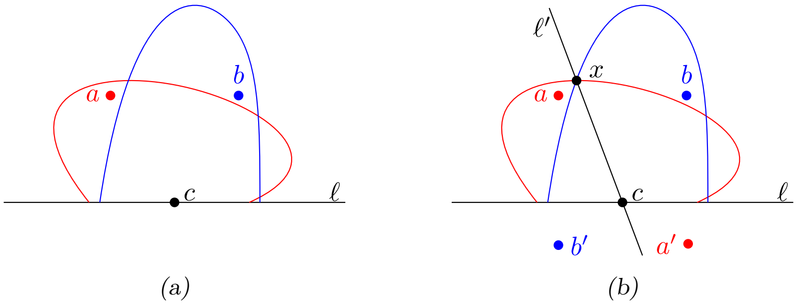

Figure 6.3. Illustration for the proof of Lemma 6.6.

Because A and B are non-homothetic, A' and B' do not coincide. Thus there exist points $a \in A' \setminus B'$ and $b \in B' \setminus A'$. We consider a line ℓ through the centroid c such that a and b lie in the same halfplane H bounded by ℓ, e.g., the line through c which is parallel to the segment ab.

Because $a \in A' \setminus B'$ and $b \in B' \setminus A'$, $c \in A' \cap B'$, the boundaries of A' and B' cross at least once in H. Let x denote the clockwise first crossing point (when turning a line in clockwise direction around c starting with ℓ) and consider the line ℓ' through x and c. Without loss of generality, we may assume that ℓ' separates a and b; otherwise there exist points with these properties and we consider them instead. Because $a \in A' \setminus B'$ and c is the centroid of both sets, there exists a point $b' \in B'$ in the halfplane bounded by ℓ' not containing b. Similarly, we obtain a point $a' \in A' \setminus B'$ in the halfplane bounded by ℓ' not containing a. Note that the points appear in the cyclic order a, b, a', b' around c. By convexity of A' and B', the points a, b, a', b' lie in convex position. Consequently, $A' \setminus B'$ and $B' \setminus A'$ cannot be separated by a line. $\square$

With the help of Lemma 6.6, we show the following.

Lemma 6.7. — *Let $C \subset \mathbb{R}^3$ be an H-difference-separable convex compact set. Then for each vector v, all non-empty sections of C orthogonal to v are homothets of each other.*

Proof. — Suppose there exist two parallel sections A and B of C that are not homothetic. Note that we may assume that these sections contain interior points of C; otherwise we slightly translate the supporting planes of the sections. Then we translate a copy C' of C such that $A = C \cap H$ and $B = C' \cap H$ for some hyperplane H orthogonal to v. By Lemma 6.6, we can scale C' around a center within H and move it orthogonally to v, such that for the resulting copy C'' and $B'' := C'' \cap H$, the sets $B'' \setminus A$ and $A \setminus B''$ cannot be separated by a line. Therefore any plane separating $C \setminus C''$ and $C'' \setminus C$ must be equal to H. However, because A and B are interior sections of C, and C is compact and convex, $C'' \setminus C$ and $C \setminus C''$ both occur on both sides of H: Let x be point of $C \cap C''$ on either side of H and let y be a point in $A \setminus B''$ (or $B'' \setminus A$). By convexity the segment xy is contained in C (or C'') and by compactness of C'' (or C) some point strictly between x and y does not belong to C'' (or C), see also Figure 6.4. Therefore, H does not separate $C'' \setminus C$ and $C \setminus C''$ and hence C is not H-difference-separable, a contradiction. $\square$

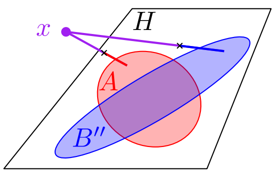

Figure 6.4. Illustration for the proof of Lemma 6.7.

Together Lemma 6.7 and Theorem 6.5 yield Theorem 1.7.

7. Future Directions

We conclude with a list of interesting open problems: We have seen that for bounded sets T-difference-separability and H-difference-separability coincide in two dimensions. Does this remain true in higher dimensions? In other words, is T-difference-separability a sufficient condition for H-difference-separability? Moreover, we are not aware of interesting compact T-difference-separable (or H-difference-separable) sets in higher dimensions beyond ellipsoids. Are these families equivalent to ellipsoids?

Its natural to wonder whether any of the conditions of Theorems 1.3 and 1.4 can be relaxed. Firstly, it seems plausible that the separability conditions can be weakened, however a proof of $\exists\mathbb{R}$-hardness would most likely require new techniques. Secondly, the NP-membership of recognition for families of homothets of a given polygon show the need for *some* curvature in order to show $\exists\mathbb{R}$-hardness. We wonder if it is sufficient for $\exists\mathbb{R}$-hardness to assume curvature at only one boundary part instead of two opposite ones. An interesting starting point could be the family of semidisks. Another open question is to consider families that include rotated copies of a fixed geometric object. Allowing for rotation, it is conceivable that $\exists\mathbb{R}$-hardness even holds for polygons.

BIBLIOGRAPHY

[1] M. Abrahamsen, Covering Polygons is Even Harder, in *IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS 2022)*, 2022, p. 375–386, https://doi.org/10.1109/FOCS52979.2021.00045.

[2] M. Abrahamsen, A. Adamaszek & T. Miltzow, The Art Gallery Problem is $\exists\mathbb{R}$-complete, in *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2018)*, 2018, p. 65–73, https://doi.org/10.1145/3188745.3188868.

[3] M. Abrahamsen, L. Kleist & T. Miltzow, Training Neural Networks is $\exists\mathbb{R}$-complete, in *Advances in Neural Information Processing Systems (NeurIPS 2021)*, 2021, p. 18293–18306, https://proceedings.neurips.cc/paper_files/paper/2021/file/9813b270ed0288e7c0388f0fd4ec68f5-paper.pdf.

[4] ———, Geometric Embeddability of Complexes is $\exists\mathbb{R}$-complete, To appear in: Symposium on Computational Geometry (SOCG 2023), 2023, https://arxiv.org/abs/2108.02585.

[5] M. Abrahamsen, T. Miltzow & N. Seiferth, Framework for $\exists\mathbb{R}$-Completeness of Two-Dimensional Packing Problems, in *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS 2020)*, 2020, p. 1014–1021, https://doi.org/10.1109/FOCS46700.2020.00098.

[6] B. Aronov, E. Ezra & M. Sharir, Small-Size ϵ-Nets for Axis-Parallel Rectangles and Boxes, *SIAM Journal on Computing (SiComp)* **39** (2010), no. 7, p. 3248–3282, https://doi.org/10.1137/090762968.

[7] M. Axenovich & T. Ueckerdt, Density of range capturing hypergraphs, *Journal of Computational Geometry (JoCG)* **7** (2016), no. 1, https://doi.org/10.20382/jocg.v7i1a1.

[8] M. L. T. Berthelsen & K. A. Hansen, On the Computational Complexity of Problems About Multi-player Nash Equilibria, in *International Symposium on Algorithmic Game Theory*, LNCS, vol. 11801, 2019, p. 153–167, https://doi.org/10.1007/978-3-030-30473-7_11.

[9] D. Bertschinger, N. El Maalouly, L. Kleist, T. Miltzow & S. Weber, The Complexity of Recognizing Geometric Hypergraphs, in *Graph Drawing and Network Visualization (GD'23)*, 2023, p. 163–179, https://doi.org/10.1007/978-3-031-49272-3_12.

[10] D. Bertschinger, C. Hertrich, P. Jungeblut, T. Miltzow & S. Weber, Training Fully Connected Neural Networks is $\exists\mathbb{R}$-Complete, 2022, https://arxiv.org/abs/2204.01368.

[11] V. Bilò & M. Mavronicolas, A Catalog of $\exists\mathbb{R}$-Complete Problems About Nash Equilibria in Multi-Player Games, in *33rd Symposium on Theoretical Aspects of Computer Science (STACS 2016)* (LIPIcs), 2016, p. 17:1–17:13, https://doi.org/10.4230/LIPIcs.STACS.2016.17.

[12] ———, $\exists\mathbb{R}$-Complete Decision Problems about Symmetric Nash Equilibria in Symmetric Multi-Player Games, in *34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)* (LIPIcs), vol. 66, 2017, p. 13:1–13:14, https://doi.org/10.4230/LIPIcs.STACS.2017.13.

[13] K. S. Booth & G. S. Lueker, Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms, *Journal of Computer and System Sciences* **13** (1976), no. 3, p. 335–379, https://doi.org/10.1016/S0022-0000(76)80045-1.

[14] ———, Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms, *Journal of computer and system sciences* **13** (1976), no. 3, p. 335–379, https://doi.org/10.1016/S0022-0000(76)80045-1.

[15] M. Čadek, M. Krčál, J. Matoušek, F. Sergeraert, L. Vokřínek & U. Wagner, Computing all maps into a sphere, *Journal of the ACM* **61** (2014), no. 3, p. 1–44, https://doi.org/10.1145/2597629.

[16] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek & U. Wagner, Time computation of homotopy groups and Postnikov systems in fixed dimension, *SIAM Journal on Computing* **43** (2014), no. 5, p. 1728–1780, https://doi.org/10.1137/120899029.

[17] M. Čadek, M. Krčál & L. Vokřínek, Algorithmic solvability of the lifting-extension problem, *Discrete & Computational Geometry (DCG)* **57** (2017), no. 4, p. 915–965, https://doi.org/10.1007/s00454-016-9855-6.

[18] J. Cardinal, S. Felsner, T. Miltzow, C. Tompkins & B. Vogtenhuber, Intersection Graphs of Rays and Grounded Segments, *Journal of Graph Algorithms and Applications* **22** (2018), no. 2, p. 273–294, https://doi.org/10.7155/jgaa.00470.

[19] D. Chistikov, S. Kiefer, I. Marusic, M. Shirmohammadi & J. Worrell, On Restricted Nonnegative Matrix Factorization, in *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, LIPIcs, vol. 55, 2016, p. 103:1–103:14, https://doi.org/10.4230/LIPIcs.ICALP.2016.103.

[20] T. K. Dey & J. Pach, Extremal Problems for Geometric Hypergraphs, *Discrete & Computational Geometry* **19** (1998), no. 4, p. 473–484, https://doi.org/10.1007/PL00009365.

[21] M. G. Dobbins, A. Holmsen & T. Miltzow, A Universality Theorem for Nested Polytopes, 2019, https://arxiv.org/abs/1908.02213.

[22] M. G. Dobbins, L. Kleist, T. Miltzow & P. Rzążewski, Completeness for the Complexity Class $\forall\exists\mathbb{R}$ and Area-Universality, *Discrete & Computational Geometry (DCG)* **70** (2022), no. 1, p. 154–188, https://doi.org/10.1007/s00454-022-00381-0.

[23] J. Erickson, Optimal Curve Straightening is $\exists\mathbb{R}$-Complete, 2019, https://arxiv.org/abs/1908.09400.

[24] J. Erickson, I. van der Hoog & T. Miltzow, Smoothing the gap between NP and $\exists\mathbb{R}$, in *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS 2020)*, 2020, p. 1022–1033, https://doi.org/10.1109/FOCS46700.2020.00099.

[25] W. Evans, P. Rzążewski, N. Saeedi, C.-S. Shin & A. Wolff, Representing Graphs and Hypergraphs by Touching Polygons in 3D„ in *Graph Drawing and Network Visualization (GD 2019)*, 2019, p. 18–32, https://doi.org/10.1007/978-3-030-35802-0_2.

[26] F. Gardi, The Roberts characterization of proper and unit interval graphs, *Discrete Mathematics* **307** (2007), no. 22, p. 2906–2908, https://doi.org/10.1016/j.disc.2006.04.043.

[27] J. Garg, R. Mehta, V. V. Vazirani & S. Yazdanbod, $\exists\mathbb{R}$-Completeness for Decision Versions of Multi-Player (Symmetric) Nash Equilibria, *ACM Transactions on Economics and Computation* **6** (2018), no. 1, article no. 1 (23 pages), https://doi.org/10.1145/3175494.

[28] M. Habib, R. McConnell, C. Paul & L. Viennot, Lex-BFS and partition refinement, with applications to transitive orientation, interval graph recognition and consecutive ones testing, *Theoretical Computer Science* **234** (2000), no. 1, p. 59–84, https://doi.org/10.1016/s0304-3975(97)00241-7.

[29] D. Haussler & E. Welzl, ε-nets and simplex range queries, *Discrete & Computational Geometry* (1987), p. 127–151, https://doi.org/10.1007/BF02187876.

[30] M. Hoffmann, T. Miltzow, S. Weber & L. Wulf, Recognition of Unit Segment Graphs is $\exists\mathbb{R}$-complete, Unpublished, in preparation.

[31] P. Jungeblut, L. Kleist & T. Miltzow, The Complexity of the Hausdorff Distance, in *38th International Symposium on Computational Geometry (SoCG 2022)*, LIPIcs, vol. 224, 2022, p. 48:1–48:17, https://doi.org/10.4230/LIPIcs.SoCG.2022.48.

[32] S. Kakeya, On some properties of convex curves and surfaces, *Tohoku Mathematical Journal, First Series* **8** (1915), p. 218–221, https://www.jstage.jst.go.jp/article/tmj1911/8/0/8_0_218/_pdf/-char/ja.

[33] R. J. Kang & T. Müller, Sphere and Dot Product Representations of Graphs, *Discrete & Computational Geometry* **47** (2012), no. 3, p. 548–568, https://doi.org/10.1007/s00454-012-9394-8.

[34] J. Kratochvíl & J. Matoušek, Intersection Graphs of Segments, *Journal of Combinatorial Theory, Series B* **62** (1994), no. 2, p. 289–315, https://doi.org/10.1006/jctb.1994.1071.

[35] P. J. Looges & S. Olariu, Optimal greedy algorithms for indifference graphs, *Computers & Mathematics with Applications* **25** (1993), no. 7, p. 15–25, https://doi.org/10.1016/0898-1221(93)90308-I.

[36] A. Lubiw, T. Miltzow & D. Mondal, The Complexity of Drawing a Graph in a Polygonal Region, in *Graph Drawing and Network Visualization (GD 2018)*, LNCS, vol. 11282, 2018, p. 387–401, https://doi.org/10.1007/978-3-030-04414-5_28.

[37] L. Ma, Bisectors and Voronoi Diagrams for Convex Distance Functions, PhD Thesis, FernUniversität Hagen, 2000.

[38] J. Matoušek, E. Sedgwick, M. Tancer & U. Wagner, Embeddability in the 3-sphere is decidable, *Journal of the ACM* **65** (2018), no. 1, p. 1–49, https://doi.org/10.1145/2582112.2582137.

[39] J. Matoušek, R. Seidel & E. Welzl, How to Net a Lot with Little: Small ϵ-Nets for Disks and Halfspaces, in *Sixth Annual Symposium on Computational Geometry (SoCG 1990)*, 1990, p. 16–22, https://doi.org/10.1145/98524.98530.

[40] J. Matoušek, M. Tancer & U. Wagner, Hardness of embedding simplicial complexes in $\mathbb{R}^d$, *JEMS* **13** (2011), no. 2, p. 259–295, https://doi.org/10.4171/JEMS/252.

[41] C. McDiarmid & T. Müller, Integer realizations of disk and segment graphs, *Journal of Combinatorial Theory, Series B* **103** (2013), no. 1, p. 114–143, https://doi.org/10.1016/j.jctb.2012.09.004.

[42] A. d. Mesmay, Y. Rieck, E. Sedgwick & M. Tancer, Embeddability in $\mathbb{R}^3$ is NP-hard, *Journal of the ACM* **67** (2020), no. 4, article no. 20 (29 pages), https://doi.org/10.1145/3396593.

[43] T. Miltzow & R. F. Schmiermann, On Classifying Continuous Constraint Satisfaction Problems, in *IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS 2021)*, 2022, p. 781–791, https://doi.org/10.1109/FOCS52979.2021.00081.

[44] N. E. Mnëv, The Universality Theorems on the Classification Problem of Configuration Varieties and Convex Polytopes Varieties, in *Topology and Geometry — Rohlin Seminar*, LNM, vol. 1346, Springer, 1988, p. 527–543, https://doi.org/10.1007/BFb0082792.

[45] S. Nakagawa, On some theorems regarding ellipsoid, *Tohoku Mathematical Journal, First Series* **8** (1915), p. 11–13, https://www.jstage.jst.go.jp/article/tmj1911/8/0/8_0_11/_pdf/-char/ja.

[46] J. Pach & G. Tardos, TIGHT LOWER BOUNDS FOR THE SIZE OF EPSILON-NETS, *Journal of the American Mathematical Society* **26** (2013), no. 3, p. 645–658, https://doi.org/10.1090/S0894-0347-2012-00759-0.

[47] J. Pach & G. Woeginger, Some New Bounds for Epsilon-Nets, in *Sixth Annual Symposium on Computational Geometry (SoCG 1990)*, 1990, p. 10–15, https://doi.org/10.1145/98524.98529.

[48] D. Rautenbach & J. L. Szwarcfiter, Unit and single point interval graphs, *Discrete Applied Mathematics* **160** (2012), no. 10-11, p. 1601–1609, https://doi.org/10.1016/j.dam.2012.02.014.

[49] J. Richter-Gebert & G. M. Ziegler, Realization Spaces of 4-Polytopes are Universal, *Bulletin of the American Mathematical Society* **32** (1995), no. 4, p. 403–412, https://doi.org/10.1090/S0273-0979-1995-00604-X.

[50] M. Schaefer, Complexity of Some Geometric and Topological Problems, in *Graph Drawing*, Springer, 2010, p. 334–344, https://doi.org/10.1007/978-3-642-11805-0.

[51] ———, Realizability of Graphs and Linkages, Thirty Essays on Geometric Graph Theory, p. 461–482, Thirty Essays on Geometric Graph Theory, Springer, 2013.

[52] ———, Complexity of Geometric k-Planarity for Fixed k, *Journal of Graph Algorithms and Applications* **25** (2021), no. 1, p. 29–41, https://doi.org/10.7155/jgaa.00548.

[53] M. Schaefer, J. Cardinal & T. Miltzow, The existential theory of the reals as a complexity class: A compendium, 2024, https://arxiv.org/abs/2407.18006.

[54] M. Schaefer, E. Sedgwick & D. Štefankovič, Recognizing string graphs in NP, *Journal of Computer and System Sciences* **67** (2003), no. 2, p. 365–380, https://doi.org/10.1016/S0022-0000(03)00045-X.

[55] M. Schaefer & D. Štefankovič, Fixed Points, Nash Equilibria, and the Existential Theory of the Reals, *Theory of Computing Systems* **60** (2017), p. 172–193, https://doi.org/10.1007/s00224-015-9662-0.

[56] ———, The Complexity of Tensor Rank, *Theory of Computing Systems* **62** (2018), no. 5, p. 1161–1174, https://doi.org/10.1007/s00224-017-9800-y.

[57] Y. Shitov, A Universality Theorem for Nonnegative Matrix Factorizations, 2016, https://arxiv.org/abs/1606.09068.

[58] ———, The Complexity of Positive Semidefinite Matrix Factorization, *SIAM Journal on Optimization* **27** (2017), no. 3, p. 1898–1909, https://doi.org/10.1137/16M1080616.

[59] P. W. Shor, Stretchability of Pseudolines is NP-Hard., in *Applied Geometry And Discrete Mathematics* (P. Gritzmann & B. Sturmfels, eds.), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 4, 1991, p. 531–554, https://doi.org/10.1090/dimacs/004/41.

[60] A. Skopenkov, Extendability of simplicial maps is undecidable, 2020, https://arxiv.org/abs/2008.00492.

[61] S. Smorodinsky, On the chromatic number of geometric hypergraphs, *SIAM Journal of Discrete Mathematics* **21** (2007), no. 3, p. 676–687, https://doi.org/10.1137/050642368.

[62] J. Spinrad, Recognition of circle graphs, *Journal of Algorithms* **16** (1994), no. 2, p. 264–282, https://doi.org/10.1006/jagm.1994.1012.

[63] J. Stade, Complexity of the boundary-guarding art gallery problem, 2022, https://arxiv.org/abs/2210.12817.

[64] P. J. Tanenbaum, M. T. Goodrich & E. R. Scheinerman, Characterization and recognition of point-halfspace and related orders, in *Graph Drawing*, Springer, 1995, p. 234–245, https://doi.org/10.1007/3-540-58950-3_375.

[65] L. Tuncel, S. Vavasis & J. Xu, Computational complexity of decomposing a symmetric matrix as a sum of positive semidefinite and diagonal matrices, 2022, https://arxiv.org/abs/2209.05678.

Manuscript received 3rd May 2024,
accepted 27th January 2025.

Daniel BERTSCHINGER
ETH Zürich
Department of Computer Science
Switzerland

Nicolas EL MAALOULY
ETH Zürich
Department of Computer Science
Switzerland

Linda KLEIST
Universität Potsdam
Institute of Computer Science
Germany
kleist@cs.uni-potsdam.de

Tillmann MILTZOW
Utrecht University
Department of Information and Computing Sciences
the Netherlands
t.miltzow@uu.nl

Simon WEBER
ETH Zürich
Department of Computer Science
Switzerland
simon.weber@inf.ethz.ch

Innovations in Graph Theory is a member of the Mersenne Center for Open Publishing
ISSN: 3050-743X